AMERICAN JOURNAL OF INDUSTRIAL MEDICINE 39:477–486 (2001)

An Application of Hierarchical Regression in the Investigation of Multiple Paternal Occupational Exposures and Neuroblastoma in Offspring

Anneclaire J. De Roos, MPH, PhD,[1]* Charles Poole, ScD,[2] Kay Teschke, PhD,[3] and Andrew F. Olshan, PhD[2]

Background *We used hierarchical regression to study the effects of 46 paternal occupational exposures on the incidence of neuroblastoma in offspring.*
Methods *The study population included 405 cases and 302 controls. The effect of each exposure was estimated using both conventional maximum likelihood and hierarchical regression.*
Results *Using hierarchical regression, overall precision was greatly enhanced compared to the conventional analysis. In addition, adjustment of effect estimates based on prespecified prior distributions of the true effect parameters allowed a more consistent interpretation across the entire panel of exposures. Estimates for several metals and solvents were shrunk close to the null value, whereas estimates for several thinner solvents, diesel fuel, solders, wood dust, and grain dust remained moderately elevated.*
Conclusions *Hierarchical regression may mitigate some of the problems of the conventional approach by controlling for correlated exposures, enhancing the precision of estimates, and providing some adjustment of estimates based on prior knowledge.* Am. J. Ind. Med. 39:477–486, 2001. *© 2001 Wiley-Liss, Inc.*

KEY WORDS: hierarchical regression; occupation; neuroblastoma; childhood cancer; multiple comparisons

INTRODUCTION

When studying associations between rare diseases and occupational exposures, the population-based case-control study is the most practical design. In conducting these studies, information is usually collected by questionnaire on numerous occupational exposures, in order to most efficiently 'screen' a long list of chemical and physical agents. Effect estimates are traditionally calculated by including each exposure in a separate logistic regression model, along with any potential confounders. There are several problems with this approach. First, people often experience combined exposures to different agents in the

[1]Occupational Epidemiology Branch, National Cancer Institute, National Institutes of Health, Rockville, Maryland.
[2]Department of Epidemiology, School of Public Health, University of North Carolina, Chapel Hill, North Carolina.
[3]Department of Health Care Epidemiology, University of British Columbia, Vancouver, BC, Canada.
Institution at which the work was performed: University of North Carolina at Chapel Hill.
Contract grant sponsor: National Institutes of Health; Contract grant number: CA57004 and T32-ES-07018.
*Correspondence to: Anneclaire J. De Roos, Occupational Epidemiology Branch, National Cancer Institute, 6120 Executive Boulevard, EPS 8091, MSC 7240, Bethesda, MD 20892-7240. E-mail: deroosa@mail.nih.gov

Accepted 21 January 2001

workplace, and the conventional analysis in which each exposure is included in a separate model does not account for probable correlation between occupational exposures. Second, the precision of estimates tends to vary considerably among the different exposures, depending on the number of persons exposed to each agent. Third, the occurrence of false associations by chance is of concern as in any epidemiologic study. These problems are particularly troubling in the situation in which multiple exposures are evaluated based on little prior knowledge. Presumably, results from such a preliminary study would be followed by research focusing on a few suspect agents, with epidemiologic studies using more sophisticated exposure assessment methods, or laboratory studies to determine effects of suspect agents on in vitro or animal models for the disease in question. However, when multiple effect estimates are elevated, the imprecision of some estimates in addition to potential confounding or generation of false positive associations, makes interpretation of the entire panel of results difficult. It would be helpful to minimize the total error in such case-control analyses to clarify a focus for further research.

Hierarchical regression, also known as multilevel or random-coefficient modeling, is a statistical method that can greatly improve the accuracy of unstable estimates, especially when studying effects of multiple exposures with limited data [Greenland, 1994, 2000, 2001]. In this type of analysis, disease outcomes are regressed on multiple exposures in a first-stage model. The beta coefficients from the first stage are then modeled as values of the outcome variable in a second-stage linear regression model, as a function of second-stage or 'prior' covariates that are thought to determine the magnitude of the true effects, or target parameters [Greenland, 1994, 2001]. Effect estimates and confidence limits are adjusted by an empirical-Bayes (EB) or semi-Bayes procedure that 'shrinks' unstable estimates toward estimated prior means of the target parameters. The shrinkage adjustments are made using the variance of an assumed prior distribution of the target parameter for each exposure. This variance is estimated in empirical-Bayes methods using an iterative procedure, or the variance can be prespecified in semi-Bayes methods by specifying a particular range in which a given proportion of the true parameter values are expected to lie.

Software for hierarchical regression modeling has not been widely available. However, a procedure written in SAS/IML for conducting multi-stage modeling of multiple exposures has recently been posted on the worldwide web (http://darwin.cwru.edu/∼witte/episoft.html) [Witte et al., 1998], and SAS Proc GLIMMIX can also be adapted to this purpose [Witte et al., 2001]. We used these methods in our recent study of the effects of paternal occupational exposures on the incidence of neuroblastoma, a childhood cancer, in offspring. In this study, exposures to 46 specific chemical and physical agents were examined. We used the SAS/IML procedure to conduct hierarchical regression using semi-Bayes and empirical-Bayes methods to generate adjusted odds ratios and confidence limits for the effects of paternal occupational exposures on the incidence of neuroblastoma. Results from hierarchical regression models generated by specifying different prior distributions were compared to each other, and to results from a conventional analysis.

MATERIALS AND METHODS

Study population. The study population for this case-control study of neuroblastoma is described in detail elsewhere [Olshan et al., 1999]. In brief, cases were patients under the age of 19 years with a confirmed new diagnosis of neuroblastoma between 1 May 1992 and 30 April 1994, registered at any of 139 participating hospitals in the United States and English-speaking Canada. The hospitals were members of either of two pediatric collaborative clinical trials groups, the Children's Cancer Group or Pediatric Oncology Group. Of the families contacted, we enrolled 538 cases (73% of those eligible). One matched control for each of 504 cases was selected by random-digit telephone dialing (RDD). The response proportion for the RDD screening was 74%. Controls were individually caliper-matched to cases on date of birth (± 6 months for cases ≤ 3 years of age, ± 1 year for cases > 3 years of age).

Data collection. A telephone interview was conducted with each mother and with the father when available. The interview included questions on demographic characteristics such as parental age, race, and education. Occupational history was obtained, including information on dates of employment, names of employers, occupations, industries, job titles, specific duties, and hours per week. For each job held during the 2-year period prior to the child's date of birth, fathers were asked if they had been exposed to electrical equipment or radiation sources, chemicals, dusts, fumes, gases, vapors, or oils. Occupational exposure information was available for a total of 707 fathers (405 case fathers, 302 control fathers).

First-stage exposures. In this study, specific chemical and physical agents were the first-stage exposures of primary interest. A review of self-reported occupational exposures was conducted by an industrial hygienist (IH) (K. Teschke), to increase the specificity of exposure variables by reducing the number of false positives in the group classified as exposed to each agent. The review was blinded to case or control parent status. The IH review covered all reported information for each exposure including occupation, industry, hours of exposure per week, form of the substance, route of exposure, use of protective equipment or clothing, work activities, and average distance from electrical equipment. A father was coded as exposed to a chemical

substance or compound if the IH review determined 'probable' exposure to that agent in any job. A father was coded as exposed to electromagnetic fields or radiation (ionizing radiation, radiofrequency fields, or extremely low frequency fields) if the IH review determined 'probable' exposure, in any job, to equipment that produces high levels of one of these frequencies. Each occupational exposure was coded as an indicator variable (1 = exposed, 0 = unexposed). In total, 46 paternal occupational exposures were coded and analyzed.

First-stage covariates. Demographic characteristics thought to be potential confounders of associations between occupational exposures and neuroblastoma were maternal education, maternal race, and maternal age at birth of the index child. The variables were coded using indicator variables for the following categories: maternal education (less than high school graduate, high school graduate and/or some college, college degree or more as the referent), maternal race (white as the referent, black, Hispanic, other), and maternal age at birth of the index child (<18, 18–39 years as the referent, ≥ 40 years).

In order to retain information on all 405 case fathers and 302 control fathers in the analysis of occupational exposures, we decided to conduct unmatched analyses by unconditional logistic regression with adjustment for the matching factor using covariates. The matching factor, child's age, was coded as a set of indicator variables with strata as fine as numbers would allow (6-month intervals for ages ≤ 3 years; 2-year intervals for ages 3–11 years, one variable for ages >11 years).

Second-stage covariates. The second-stage matrix contained variables thought to determine the magnitude of, or explain some of the variability between, the individual target parameters. These second-stage, or prior, covariates were indicator variables representing subsets of target parameters within which the parameters were regarded as 'exchangeable', or as draws from a common prior distribution [Greenland, 1994, 2000]. We defined exchangeable categories by grouping the first-stage occupational exposures according to similarities in physicochemical properties, as halogenated hydrocarbons (HCs), nonvolatile HCs, volatile HCs, metals, paints, thinner solvents, wood-derived substances, grain-derived dusts, and non-ionizing electromagnetic field exposures (see Table I). Some exposures were included in more than one category (e.g., oil-based paints, turpentine). Other individual exposures did not have physicochemical properties that we would expect to translate into exchangeable biologic effects for neuroblastoma; these exposures were therefore not grouped together (e.g., as in groups of pesticides or dusts). There is little prior information on the potential for any of the agents to act as transgenerational carcinogens through a paternally mediated mechanism; thus, no second-stage covariates indicating toxicity were used.

TABLE I. Coding of Second-Stage Covariates: Categories of Exchangeability Between True Effect Parameters for Occupational Exposures[a,b]

Halogenated hydrocarbons	Metals
Carbon tetrachloride	Brass
Chloroform	Bronze
Freon	Galvanized steel
Methylene chloride	High-speed steel
Perchloroethylene	Mild steel
Trichloroethylene	Stainless steel
	Alloys (NOS)
Non-volatile hydrocarbons	Metals (NOS)
Cutting oil	Solders (NOS)
Diesel fuel	
Kerosene	Thinner solvents
Lubricating oil	Lacquer thinner
	Mineral spirits
Volatile hydrocarbons	Oil-based paints
Acetone	Paint thinner
Alcohols	Turpentine
Benzene	
Gasoline	Wood-derived substances
Glycols or glycol ethers	Turpentine
Lacquer thinner	Wood dust
Methyl ethyl ketone	
Mineral spirits	Grain-derived dusts
Naphtha	Flour dust
Paint thinner	Grain dust
Toluene	
Turpentine	Non-ionizing EMF
White gas	Extremely low frequency fields
Xylene	Radiofrequency fields
Paints	
Oil-based paints	
Water-based paints	

[a]Other chemicals analyzed but not included in any grouping: plastics, synthetics, or resins (NOS); cardboard dust; rubber dust; herbicides (NOS), insecticides (NOS), ionizing radiation.
[b]NOS, not otherwise specified; EMF, electric and magnetic fields.

The second-stage, or Z-matrix, was structured with one row for each of the occupational exposures, j. Each row was composed of 10 elements; column one contained a value of '1' denoting the presence of an intercept, and the following nine columns contained values for the each of the prior covariates, z_{ij}; a '1' if the exposure was present in that category and a '0' if not. An exposure that appeared in more than one category thus had multiple '1's in its row of the Z-matrix.

Statistical analyses. In our conventional analysis, each occupational exposure was evaluated in a separate unconditional logistic regression model, along with indicator variables representing child's age (the matching factor),

and the demographic covariates. Exposure odds ratios estimating incidence rate ratios were estimated using maximum likelihood.

In the first-stage model of the hierarchical regression analyses, neuroblastoma disease status was regressed simultaneously on the 46 paternal occupational exposures, child's age, and the demographic covariates. This model took the form: $\Pr(y = 1| x, w) = \text{expit}(\alpha + X\beta + W_\gamma)$, where X represents an n-row matrix of occupational exposures and W represents an n-row matrix of potential confounders, where n represents the number of subjects in the study, and expit (●) is the logistic function exp (●)/(1 + exp (●)) [Greenland, 1998].

The estimated beta coefficients for the 46 occupational exposures in the first-stage model were then regressed in a second-stage linear regression model as a function of the prior covariates. The second-stage model should incorporate what is known about each target parameter, β_j, prior to seeing the study data [Greenland, 1994, 2001]. Therefore, a prior 'distribution' was defined for the true effect parameter for each occupational exposure, with a prior mean dependent on the joint distribution of second-stage covariates, and a prespecified prior variance for each parameter. For each occupational exposure, j, the second-stage model took the form: $\beta_j = \pi_1 z_{1j} + \pi_2 z_{2j} + \pi_3 z_{3j} + \pi_4 z_{4j} + \pi_5 z_{5j} + \pi_6 z_{6j} + \pi_7 z_{7j} + \pi_8 z_{8j} + \pi_9 z_{9j} + \delta_j = z_j\pi + \delta_j$, where π is the column vector containing the second-stage parameters, π_1 through π_9, and δ_j the deviation of the effect of occupational exposure, j, from the sum $z_j\pi$. Based on the prior distributions, the second-stage model assumed that each target parameter deviates randomly around the linear term on the right hand side of the equation [Greenland, 1992, 1993, 1994, 1998]. In other words, target parameters for occupational exposures with the same values for the second-stage covariates were assumed to have been randomly sampled from a common underlying distribution, with an unknown mean. In addition to hierarchical models using the prior covariates to determine prior means of the target parameters, we ran one hierarchical model with an intercept-only second-stage matrix, containing no prior covariates. Because our prior covariates were crudely specified categories of exchangeability, we wanted to compare a hierarchical model using no prior covariates to one using our crudely specified prior covariates, to assess the benefit of such a procedure in the face of little or no prior information. In this intercept-only model, all the target parameters were assumed to have been sampled from a common distribution with an unknown mean.

The deviation δ_j is called the residual effect of occupational exposure j; it represents effects of exposure j that are not captured by the sum $z_j\pi$ [Greenland, 1994, 1998, 2001; Witte et al., 1994], or effects above and beyond those accounted for by the 'group' effects of the second-stage covariates. Residual effects can arise from information not included in the Z-matrix, such as other information that could potentially explain variability between the target parameters; for example, detailed measures of genotoxicity and carcinogenicity. These residual effects δ_j are usually assumed to be independent random quantities having means of zero and variance τ^2, where τ^2 may be fixed in advance using background information, as in semi-Bayes analyses, or may be estimated from the data, as in empirical-Bayes analyses [Greenland, 1992, 1993, 1994, 1998, 2000, 2001; Greenland and Poole, 1994; Witte et al., 1994, 1998, 2001]. A large value for τ^2 would imply that there are likely to be substantial effects of an exposure beyond those explained by the second-stage covariates, whereas a small value for τ^2 would translate into a relatively tight range for the residual effects, and would imply a greater confidence that the effects of the exposure act through mechanisms that are almost completely mediated through the second-stage covariates. In our semi-Bayes analyses, we used different prespecified values for τ^2 in different hierarchical models to observe the sensitivity of our results to the choice. Although the hierarchical model can be generalized by allowing τ^2 to vary for different first-stage exposures [Greenland, 1994], we did not feel that we had better prior information for any specific exposure compared to the others; therefore, in each analysis, the same value for τ^2 was assigned to all exposures. Because our Z-matrix was rather crudely defined, we started with a liberal prespecified range for the residual effects of the occupational exposures. We assumed, with 95% certainty, that the rate ratio for each occupational exposure, after adjusting for the second-stage covariates, would fall within a 10-fold range (e.g., between 0.5 and 5.0), or that the δ_j is, with 95% certainty, an interval 2.3 units wide on the log rate ratio scale (e.g., $\ln(0.5) = -0.7$ and $\ln(5.0) = 1.6$). Thus, assuming that the δ_j are normally distributed with standard deviation of τ, $(1.96)(2)\tau = 2.3$, or $\tau = 0.59$, and the prior residual variance, τ^2, is 0.35. We also conducted two other hierarchical regression analyses using prespecified five-fold and 2.5-fold 95% ranges to estimate the prior variance, to see if the results were sensitive to the choice [Greenland, 1992, 1993, 1994, 2001; Greenland and Poole, 1994; Witte et al., 1994]. In addition, we performed a hierarchical regression analysis using empirical-Bayes estimation of the residual variance.

The first- and second-stage models together constitute a two-stage hierarchical regression model [Greenland, 1994, 1998, 2000, 2001]. The prior mean for each occupational exposure was estimated by substituting the estimated π_1 through π_9 into the equation, $\hat{\beta} = Z_j\hat{\pi}$. An average of the estimated prior mean vector and the vector of maximum likelihood estimates from the first-stage model, weighted by the covariance matrix of the first-stage estimate and τ^2, respectively, gives the estimated posterior coefficient for each exposure [Greenland, 1992, 1993, 1994; Witte and Greenland, 1996].

RESULTS

The effects of paternal occupational exposures on the incidence of neuroblastoma in offspring, estimated from the various analyses, are shown in Table II. The results from our conventional analysis using maximum likelihood estimation for each agent separately are presented as Analysis 1. Effect estimates for some exposures are relatively precise (e.g., wood dust, 95% CL ratio = 3.5 [CL ratio = upper confidence limit divided by lower confidence limit]), while others are very imprecise (e.g., turpentine, CL ratio = 18.7). The precision of these estimates is primarily dependent on the number and distribution of exposed cases and controls. While some imprecise estimates are sufficiently elevated to pique interest for further study (e.g., turpentine, OR = 10.4; CL ratio = 18.7), other imprecise estimates are only moderately elevated (e.g., high-speed steel, OR = 2.0; CL ratio = 15.4), making prioritization of future research difficult. Although each estimate can be interpreted in a straightforward manner as the observed association between the exposure and the disease outcome, some of the associations may result from confounding by combined exposures to different agents in the workplace. As in any statistical analysis, false positive associations may also have occurred by chance. Analysis 1 can be conceived of as a special case of hierarchical modeling, in which the prior residual variance is set at infinity (i.e., odds ratios of any magnitude are a priori equally likely to occur), and the true effect parameter for each occupational exposure has its own independent prior mean [Greenland, 1992]. Therefore, a second-stage model defined in this way would do nothing to adjust the estimates or confidence limits from the first stage.

The results from the first-stage model of the hierarchical regression, estimated by simultaneously modeling all occupational exposures, child's age, and the demographic covariates using maximum likelihood estimation, are presented as Analysis 2. Many of these effect estimates are wildly imprecise (e.g., turpentine, CL ratio = 54). Although the advantage of such an analysis is to control for the correlation between exposures that occur together in the workplace, the imprecision resulting from the inclusion of so many variables in one model makes the estimates virtually meaningless. This type of model is really only useful as a first-stage model for the hierarchical regression.

Analysis 3 shows the results from hierarchical regression modeling using an intercept-only second-stage matrix, with a prespecified 10-fold range for the true residual effect parameters given 95% certainty. There are no prior covariates in the second-stage equation for this analysis. The results from Analysis 3 show greatly increased precision of estimates that were previously unstable (e.g., turpentine, CL ratio = 6.5), and the effect estimates for the previously unstable estimates have been shrunk closer toward the mean of all the estimates. Analysis 3 provides somewhat of a remedy for each of the three problems of the conventional analysis (Analysis 1). Inclusion of all occupational exposures in a single model controls for potential confounding by combined occupational exposures. The enhanced precision of posterior estimates in spite of sparse data gives us greater confidence in interpreting the observed results. In addition, the shrinkage of estimates toward the mean of all estimates, within an a priori probable range for the parameters, theoretically provides some adjustment for results that may occur by chance. Any outlying association will be penalized in such a way that its posterior estimate will be closer to those of the other occupational exposures. Therefore, if the average effect of all occupational exposures is the null value, any non-null association will be somewhat attenuated. The intercept-only second-stage model implies dependency among the entire group of occupational exposures; undoubtedly, the adjustment procedure could be much improved if it was based on the distribution of prior covariates that are thought to explain some of the variability between the effects of different occupational exposures.

By adding a Z-matrix containing prior covariates to the second-stage model of the hierarchical regression analysis, we force the estimates for first-stage exposures with the same values of second-stage covariates to be more similar to each other than to those with other values. However, the results for Analysis 4 show that in this situation, with little prior information, the grouping of exposures into categories of exchangeability based on physicochemical properties adds little information to the overall analysis. With a few exceptions, the magnitude and precision of estimates from Analysis 4 are quite similar to those from Analysis 3, in which no prior covariates were added to the second stage. Exceptions generally occur where an exposure was included in more than one category of exchangeability. For example, turpentine was included in the categories of volatile hydrocarbons, thinner solvents, and wood-derived substances. The effect estimate is more elevated and more imprecise than that in Analysis 3 (OR = 4.8; CL ratio = 13.7), with the increased imprecision resulting from its complex joint distribution among second-stage covariates. Presumably, however, this estimate with the incorporated information from the second-stage matrix may be more accurate than the estimate from Analysis 3, albeit more imprecise. Effect estimates were not always elevated when included in a category of exchangeability with others that had elevated odds ratios; for example, the estimate for oil paint is near the null value, despite its inclusion in the thinner solvent category along with turpentine, lacquer thinner, mineral spirits, all of which had elevated odds ratios. The use of prior covariates in the Z-matrix theoretically improves the adjustment for associations distributed by random error, because the effect of each exposure is expected to be more similar to other exposures with the same values for the prior

TABLE II. Estimated Effects of Paternal Occupational Exposures (Odds Ratios and 95% Confidence Intervals) on the Incidence of Neuroblastoma in Offspring[a,b] Listed by Categories of Exchangeability[c]

Exposure	Number exposed		Analysis 1—ML[d], one stage, one model for each exposure	Analysis 2—ML, one stage, one model including all exposures	Analysis 3—HM[d], semi-Bayes, 10-fold range to estimate prior residual variance[e], no prior covariates	Analysis 4—HM, semi-Bayes, 10-fold range to estimate prior residual variance, with prior covariates	Analysis 5—HM, semi-Bayes, 5-fold range to estimate prior residual variance, with prior covariates	Analysis 6—HM, semi-Bayes, 2.5-fold range to estimate prior residual variance, with prior covariates
	case	control						
Halogenated hydrocarbons								
Carbon tetrachloride	4	4	0.6 (0.2, 2.6)	1.1 (0.1, 10.2)	0.9 (0.3, 2.5)	0.7 (0.3, 2.1)	0.7 (0.3, 1.7)	0.7 (0.4, 1.3)
Chloroform	3	2	1.2 (0.2, 7.5)	1.3 (0.1, 20.4)	1.0 (0.4, 2.9)	0.8 (0.2, 2.4)	0.7 (0.3, 1.8)	0.7 (0.4, 1.4)
Freon	9	13	0.5 (0.2, 1.1)	0.5 (0.1, 1.7)	0.6 (0.3, 1.4)	0.6 (0.2, 1.4)	0.6 (0.3, 1.3)	0.7 (0.4, 1.2)
Methylene chloride	4	4	0.7 (0.2, 2.8)	1.0 (0.2, 7.2)	0.9 (0.4, 2.4)	0.7 (0.3, 2.1)	0.7 (0.3, 1.7)	0.7 (0.4, 1.4)
Perchloroethylene	4	6	0.5 (0.1, 1.7)	0.3 (0.1, 2.7)	0.8 (0.3, 2.1)	0.6 (0.2, 1.7)	0.7 (0.3, 1.5)	0.7 (0.4, 1.3)
Trichloroethylene	9	7	0.9 (0.3, 2.5)	0.7 (0.1, 3.0)	0.9 (0.4, 2.1)	0.8 (0.3, 1.9)	0.8 (0.4, 1.6)	0.7 (0.4, 1.3)
Non-volatile hydrocarbons								
Cutting oil	16	7	1.7 (0.7, 4.2)	2.0 (0.5, 7.5)	1.2 (0.5, 2.8)	1.2 (0.5, 3.0)	1.1 (0.5, 2.4)	1.0 (0.6, 1.8)
Diesel fuel	42	21	1.5 (0.8, 2.6)	2.2 (0.9, 5.0)	1.5 (0.8, 2.9)	1.5 (0.8, 3.0)	1.3 (0.7, 2.4)	1.1 (0.7, 1.8)
Kerosene	16	11	1.0 (0.5, 2.2)	0.4 (0.1, 1.4)	0.8 (0.4, 1.7)	0.7 (0.3, 1.7)	0.8 (0.4, 1.7)	0.9 (0.5, 1.6)
Lubricating oil or grease	56	36	1.1 (0.7, 1.8)	0.9 (0.5, 1.9)	1.0 (0.5, 1.7)	1.0 (0.5, 1.7)	1.0 (0.6, 1.7)	1.0 (0.6, 1.5)
Volatile hydrocarbons								
Acetone	23	19	0.9 (0.5, 1.7)	0.4 (0.1, 0.9)	0.6 (0.3, 1.3)	0.6 (0.3, 1.1)	0.6 (0.4, 1.2)	0.8 (0.5, 1.2)
Alcohols	35	16	1.8 (0.9, 3.3)	1.4 (0.6, 3.4)	1.3 (0.7, 2.6)	1.3 (0.7, 2.6)	1.2 (0.7, 2.1)	1.0 (0.7, 1.6)
Benzene	5	2	2.0 (0.4, 10.3)	1.4 (0.1, 15.4)	1.1 (0.4, 3.1)	1.0 (0.4, 3.0)	1.0 (0.4, 2.2)	0.9 (0.5, 1.5)
Gasoline	45	38	0.8 (0.5, 1.3)	0.4 (0.2, 0.9)	0.6 (0.3, 1.0)	0.6 (0.3, 1.0)	0.6 (0.4, 1.1)	0.7 (0.5, 1.1)
Glycols or glycol ethers	7	4	1.3 (0.4, 4.6)	1.4 (0.2, 8.3)	1.1 (0.4, 2.9)	1.1 (0.4, 2.9)	1.0 (0.4, 2.1)	0.9 (0.5, 1.5)
Lacquer thinner	36	8	3.5 (1.6, 7.8)	3.5 (1.1, 11.6)	1.9 (0.9, 4.0)	2.1 (0.9, 4.7)	1.9 (1.0, 3.7)	1.7 (1.0, 2.8)
Methyl ethyl ketone	12	6	1.4 (0.5, 3.8)	0.8 (0.2, 3.6)	1.0 (0.4, 2.3)	0.9 (0.4, 2.1)	0.9 (0.4, 1.8)	0.9 (0.5, 1.4)
Mineral spirits	26	9	2.2 (1.0, 4.9)	2.2 (0.7, 7.2)	1.4 (0.6, 3.1)	1.8 (0.8, 4.2)	1.7 (0.8, 3.5)	1.6 (0.9, 2.8)
Naphtha	6	3	1.4 (0.4, 5.9)	1.1 (0.1, 11.4)	1.1 (0.4, 2.9)	1.0 (0.4, 2.8)	0.9 (0.4, 2.1)	0.9 (0.5, 1.5)
Paint thinner	43	17	1.9 (1.0, 3.4)	0.8 (0.3, 2.1)	1.1 (0.6, 2.2)	1.2 (0.6, 2.5)	1.3 (0.7, 2.5)	1.5 (0.9, 2.4)
Toluene	10	7	1.0 (0.4, 2.7)	0.7 (0.2, 3.2)	0.9 (0.4, 2.1)	0.8 (0.3, 2.0)	0.9 (0.4, 1.7)	0.9 (0.5, 1.4)
Turpentine	25	2	10.4 (2.4, 44.8)	17.8 (2.4, 130)	2.0 (0.8, 5.2)	4.8 (1.3, 17.9)	3.8 (1.2, 12.0)	3.1 (1.2, 7.9)
White gas	5	3	1.2 (0.3, 5.3)	1.1 (0.1, 11.0)	1.1 (0.4, 3.0)	1.0 (0.3, 2.8)	0.9 (0.4, 2.1)	0.9 (0.5, 1.5)
Xylene	10	5	1.4 (0.5, 4.3)	2.3 (0.4, 14.4)	1.1 (0.4, 2.8)	1.2 (0.5, 3.0)	1.0 (0.5, 2.1)	0.9 (0.6, 1.5)

Paints								
Oil-based paints	27	14	0.9 (0.3, 3.0)	0.9 (0.3, 3.0)	1.0 (0.5, 2.0)	0.9 (0.4, 2.2)	1.0 (0.5, 2.0)	1.0 (0.5, 1.8)
Water-based paints	24	16	1.1 (0.6, 2.2)	0.6 (0.2, 1.8)	0.8 (0.4, 1.6)	0.6 (0.3, 1.3)	0.6 (0.3, 1.2)	0.6 (0.3, 1.1)
Thinner solvents								
Lacquer thinner	36	8	3.5 (1.6, 7.8)	3.5 (1.1, 11.6)	1.9 (0.9, 4.0)	2.1 (0.9, 4.7)	1.9 (1.0, 3.7)	1.7 (1.0, 2.8)
Mineral spirits	26	9	2.2 (1.0, 4.9)	2.2 (0.7, 7.2)	1.4 (0.6, 3.1)	1.8 (0.8, 4.2)	1.7 (0.8, 3.5)	1.6 (0.9, 2.8)
Oil-based paints	27	14	0.9 (0.3, 3.0)	0.9 (0.3, 3.0)	1.0 (0.5, 2.0)	0.9 (0.4, 2.2)	1.0 (0.5, 2.0)	1.0 (0.5, 1.8)
Paint thinner	43	17	1.9 (1.0, 3.4)	0.8 (0.3, 2.1)	1.1 (0.6, 2.2)	1.2 (0.6, 2.5)	1.3 (0.7, 2.5)	1.5 (0.9, 2.4)
Turpentine	25	2	10.4 (2.4, 44.8)	17.8 (2.4, 130)	2.0 (0.8, 5.2)	4.8 (1.3, 17.9)	3.8 (1.2, 12.0)	3.1 (1.2, 7.9)
Metals								
Brass	8	4	1.5 (0.4, 5.2)	0.4 (0.1, 3.9)	1.0 (0.4, 2.5)	1.0 (0.4, 2.5)	1.0 (0.5, 2.2)	1.1 (0.7, 1.8)
Bronze	4	2	1.4 (0.3, 7.9)	1.7 (0.1, 50.9)	1.0 (0.3, 2.8)	1.1 (0.4, 3.2)	1.1 (0.5, 2.4)	1.1 (0.7, 1.8)
Galvanized iron or steel	18	8	1.6 (0.7, 3.9)	1.3 (0.3, 5.0)	1.2 (0.5, 2.7)	1.1 (0.5, 2.6)	1.1 (0.6, 2.2)	1.1 (0.7, 1.8)
High-speed steel	8	3	2.0 (0.5, 7.7)	2.8 (0.3, 27.2)	1.1 (0.4, 3.1)	1.2 (0.4, 3.2)	1.1 (0.5, 2.4)	1.1 (0.7, 1.8)
Mild steel	15	7	1.6 (0.6, 4.0)	0.4 (0.1, 2.0)	1.0 (0.4, 2.3)	0.9 (0.4, 2.1)	1.0 (0.5, 1.9)	1.1 (0.7, 1.7)
Stainless steel	13	5	1.9 (0.6, 5.4)	2.8 (0.5, 16.2)	1.2 (0.5, 3.1)	1.3 (0.5, 3.3)	1.2 (0.6, 2.5)	1.1 (0.7, 1.8)
Alloys (NOS[a])	3	5	0.4 (0.9, 1.7)	0.2 (0.1, 1.5)	0.7 (0.3, 1.9)	0.7 (0.3, 2.0)	0.9 (0.4, 1.9)	1.0 (0.6, 1.7)
Metals (NOS)	10	3	2.6 (0.7, 9.5)	2.0 (0.3, 15.5)	1.2 (0.5, 3.2)	1.3 (0.5, 3.4)	1.2 (0.6, 2.5)	1.1 (0.7, 1.9)
Solders (NOS)	17	5	2.6 (0.9, 7.1)	3.4 (0.8, 14.1)	1.4 (0.6, 3.4)	1.5 (0.6, 3.6)	1.3 (0.6, 2.7)	1.2 (0.7, 1.9)
Wood-derived substances								
Turpentine	25	2	10.4 (2.4, 44.8)	17.8 (2.4, 130)	2.0 (0.8, 5.2)	4.8 (1.3, 17.9)	3.8 (1.2, 12.0)	3.1 (1.2, 7.9)
Wood dust	47	26	1.5 (0.8, 2.8)	1.8 (0.9, 3.7)	1.4 (0.8, 2.6)	1.8 (0.9, 3.5)	1.8 (0.9, 3.5)	1.8 (0.9, 3.4)
Grain-derived dusts								
Flour dust	8	6	1.3 (0.3, 5.5)	1.4 (0.3, 6.8)	1.2 (0.5, 2.9)	1.6 (0.5, 5.3)	1.6 (0.5, 5.0)	1.6 (0.6, 4.7)
Grain dust	9	6	3.2 (0.7, 15.2)	4.4 (0.7, 29.7)	1.4 (0.5, 3.7)	2.0 (0.5, 7.3)	1.8 (0.5, 6.0)	1.7 (0.6, 5.1)
Non-ionizing EMF[4]								
Extremely low frequency fields	57	34	1.2 (0.8, 1.9)	0.8 (0.5, 1.5)	1.0 (0.6, 1.6)	1.0 (0.6, 1.7)	1.0 (0.6, 1.7)	1.1 (0.7, 1.7)
Radiofrequency fields	45	27	1.3 (0.8, 2.2)	1.3 (0.7, 2.3)	1.2 (0.7, 2.0)	1.3 (0.7, 2.1)	1.2 (0.7, 2.0)	1.2 (0.8, 1.9)
Cardboard dust	14	15	1.1 (0.4, 3.0)	0.9 (0.3, 2.8)	1.0 (0.5, 2.2)	1.0 (0.4, 2.2)	1.0 (0.5, 2.0)	1.0 (0.6, 1.8)
Herbicides (NOS)	13	8	1.3 (0.5, 3.2)	1.7 (0.4, 6.4)	1.2 (0.5, 2.7)	1.2 (0.5, 2.8)	1.1 (0.6, 2.3)	1.1 (0.6, 1.8)
Insecticides (NOS)	10	7	1.1 (0.4, 3.0)	0.8 (0.2, 3.3)	1.1 (0.5, 2.4)	1.0 (0.4, 2.4)	1.0 (0.5, 2.1)	1.0 (0.6, 1.8)
Ionizing radiation	5	4	1.2 (0.3, 4.5)	1.8 (0.4, 7.4)	1.2 (0.5, 3.0)	1.2 (0.4, 3.1)	1.1 (0.5, 2.5)	1.0 (0.6, 1.9)
Plastics, synthetics, resins (NOS)	4	5	0.6 (0.2, 2.3)	0.1 (0.0, 0.9)	0.7 (0.3, 1.9)	0.7 (0.2, 2.0)	0.8 (0.3, 2.0)	0.9 (0.5, 1.8)
Rubber dust	6	5	2.8 (0.3, 25.5)	5.6 (0.4, 88.1)	1.2 (0.4, 3.5)	1.2 (0.4, 3.9)	1.1 (0.4, 2.8)	1.1 (0.6, 2.0)

[a]n, 405 Case fathers; 302 control fathers.

[b]All estimates are adjusted for child's age, maternal race, maternal age, and maternal education.

[c]Subheadings indicate second-stage 'prior' covariates or categories of exchangeability; exposures may occur in more than one category.

[d]ML, maximum likelihood estimation; HM, hierarchical model; NOS, not otherwise specified; and EMF, electric and magnetic fields.

[e]Range in which we are 95% certain that true effect parameters will lie, after adjusting for second-stage covariates.

covariates. Therefore, if one exposure in the entire category of exchangeability had a falsely elevated or attenuated association due to chance, its estimate would be forced to be more similar to the others in the group. Obviously, more sophisticated specification of the Z-matrix would provide much better information for adjustment of estimates.

Analyses 5 and 6 are similar to the hierarchical regression specified in Analysis 4, except for the smaller values of the prespecified 95% ranges to estimate the prior variances. These models demonstrate the power the investigator has in controlling how conservative the estimates should be. The extreme situation in which the prior 95% range for the true parameters is set at a 2.5-fold range, illustrates that with this tight range, the chances of observing any but the strongest associations are greatly decreased.

In the hierarchical regression using empirical-Bayes estimation of the residual variance, τ^2 was set to zero (corresponding to a 95% certainty there are no residual effects of the first-stage exposures after adjusting for the second-stage covariates), and the SAS/IML procedure generated an automatic message indicating that there was possible underdispersion and semi-Bayes methods should be considered. Because of the potential problems in the estimation of τ^2, the results from this model are not presented.

DISCUSSION

In our conventional analysis of paternal occupational exposures and the incidence of neuroblastoma in offspring, there was some indication that exposures to hydrocarbons such as diesel fuel, lacquer thinner, mineral spirits, and turpentine, metals such as stainless and high-speed steel, and dusts such as wood dust were positively associated with neuroblastoma in offspring (Analysis 1). By adjusting the estimates based on prespecified prior distributions of the true effect parameters, a more consistent interpretation of the effects across the entire panel of exposures was possible. For example, although effect estimates were originally elevated for individual metal exposures (Analysis 1), these estimates were imprecise due to small numbers of persons exposed to each agent. After adjustment in the hierarchical regression analyses, most of these effect estimates were shrunk closer to the null value (Analyses 4). In comparison to the results for volatile hydrocarbons, of which several remain moderately elevated following the shrinkage procedure, the results for metals are less convincing. Conversely, we did not originally place much emphasis on the positive association observed for grain dust, because of its extreme imprecision and modestly elevated odds ratio (Analysis 1; OR = 3.2; CL ratio = 21.7). However, the hierarchical regression results show that the effect estimate for grain dust is somewhat elevated even after the shrinkage of the imprecise effect estimate toward a prior mean (Analysis 4; OR = 2.0; CL ratio = 14.6). These findings shed some light on our original interpretation of the data. Based on the results of the hierarchical regression analyses, further research into the effects of thinner solvents (lacquer thinner, mineral spirits, and turpentine), diesel fuel, and grain-, flour and wood dusts appear warranted, although a focus on metals seems less justifiable.

Although an appropriate value for the variance of true effect parameters was uncertain, we assumed that the true rate ratio for the effect of any occupational exposure, after adjusting for the second-stage covariates, is not likely to fall outside of a 10-fold range (e.g., between 0.5 and 5.0). Our results from the semi-Bayes analyses were sensitive to this choice; the odds ratios are shrunk closer to the null value in the models with prespecified five-fold and 2.5-fold ranges, which in some cases would affect our interpretation of results. With the measurement error and misclassification inherent in occupational exposure assessment that may bias odds ratios toward the null value when errors are independent and non-differential, we felt that interpretation of results from the more liberal analysis with 10-fold range was appropriate in this type of preliminary analysis of associations. In addition, simulation studies have observed that when conducting hierarchical regression with many first-stage exposures and few second-stage covariates, overspecifying τ^2 did not harm the confidence interval coverage or the mean squared error of estimates, whereas underspecifying τ^2 did [Greenland, 1993; Witte and Greenland, 1996]. Nonetheless, the entire array of information is useful in examining the sensitivity of results to the choice of the 95% prior range for the true residual effect parameters, and in providing further confidence in some results that remain elevated when the prior residual variance is extrezmely small.

In our example, use of empirical-Bayes estimation of the prior residual variance resulted in an estimate of zero for the variance of target parameters after adjusting for the second-stage covariates. Given the crude nature of our second-stage matrix, a value of zero for the prior residual variance is definitely unrealistic, and indicates a failure of the estimation procedure. Semi-Bayes methods appear to outperform empirical-Bayes methods when the ratio of subjects to parameters is not large [Greenland, 1993], as in our study (707 subjects/46 exposures $\approx$7). Thus, semi-Bayes methods seem preferable in this type of situation in which multiple effects are being estimated in a study of limited sample size, as long as the target parameters for a set of exposures are reasonably expected to fall within a definable range.

Use of a second-stage prior model can greatly improve the accuracy of effect estimates by modeling similarities among parameters of interest [Greenland, 1992, 1993, 1994, 1997, 2000; Greenland, 2001; Greenland and Poole, 1994; Witte and Greenland, 1996; Witte et al., 1994]. However, in

our study in which there was little prior information, a crudely specified second-stage covariate matrix did little to affect most of the estimates beyond the impact of an intercept-only second-stage model. Where no reliable prior information exists, a hierarchical regression analysis with an intercept-only model and prespecified residual variance at the second stage can be useful in enhancing the precision of estimates, as in our Analysis 3. Although confidence interval coverage rates tend to decrease with decreasing numbers of second-stage covariates, one simulation study found that having an intercept-only model did not harm confidence interval coverage rates as long as the prior residual variance was not underspecified [Witte and Greenland, 1996].

Even where only crude prior information exists (as in the form of categories of exchangeability), a hierarchical model with a simplified second-stage can outperform maximum likelihood estimation in enhancing the accuracy and precision of estimates [Witte and Greenland, 1996]. Some of the estimates in our analysis were changed by the inclusion of second-stage prior covariates in the model. However, where the use of prior covariates made a difference in the magnitude of effect estimates, there was also a loss of precision. These estimates may be more accurate as a result of added information provided in the second stage; however, the loss of precision may affect interpretation. Our prior covariates were all indicator variables grouping the first-stage exposures into exchangeable categories, and stratification of some of the exposures across several categories of the prior covariates was responsible for harming the precision of their posterior estimates. In general, lower precision occurs with the inclusion of a greater number of second-stage covariates, especially when the number of first-stage exposures is large [Greenland, 1993; Witte and Greenland, 1996]. Presumably, however, with the use of a carefully specified Z-matrix, the loss of precision is offset by reduction in bias. An improvement on the categories of exchangeability in our Z-matrix would be prior covariates that describe the toxicity of occupational exposures; for example, continuous measures of genotoxic potency that would determine the effects as linear functions of the continuous prior covariates. The Z-matrix could alternatively contain components of the first-stage exposures, as in studies of the effects of individual dietary items at the first-stage, mediated through the nutrient levels contained within each dietary item at the second-stage [Witte et al., 1994, 1998, 2001].

Simulation studies of hierarchical regression [Greenland, 1993, 1997; Witte and Greenland, 1996] have not addressed validity issues such as selection bias, reporting bias, and measurement error. In a population-based case-control study such as ours, these problems may greatly affect the accuracy of effect estimates. Although we have taken extensive measures to attain complete case ascertainment, collect a representative sample of population-based controls, and reduce reporting bias by carefully reviewing each reported occupational exposure to improve the specificity of our exposure measures, biases may still exist. Some of these biases could pose particular problems unique to the results of hierarchical regression analyses. For example, in our review of self-reported occupational exposures, we noticed that the frequency of reporting was higher for identifiable substances (e.g., turpentine) and chemical products with well-known common names (e.g., gasoline), compared to less frequent reporting of individual chemicals (e.g., benzene). Because these substances occur in the same category of exchangeability in the hierarchical regression models (i.e., volatile hydrocarbons), measurement error resulting from underreporting of individual chemicals may bias results in an unknown direction, not only for the substance with poor reporting, but for all other exposures in the category, since posterior estimates have been shrunk toward the common prior mean. Nothing conclusive is known about the accuracy of these self-reported occupational exposures; therefore, it is impossible to know to what extent results may have been influenced by differential reporting of the individual exposures. As in every statistical analysis, the quality of data will in part determine the quality of results.

With the development and increasing availability of software to perform hierarchical regression analyses, these procedures are becoming more accessible to application in epidemiology. The SAS/IML procedure we used is simple, and solely requires structuring as matrices in the SAS/IML language the results from the first-stage model, the prior residual variance, and second-stage covariates. However, the printed version [Witte et al., 1998] has errors which are corrected in the downloadable version posted on the Web. Also the procedure uses a weighted least-squares algorithm that tends to produce excessively wide intervals in small samples [Greenland, 1993]. The more sophisticated penalized-likelihood algorithm used by SAS GLIMMIX does not suffer from this problem [Greenland, 1997] and is also simple to use for epidemiologic analysis [Witte et al., 2001].

With relative ease of use, such analyses should become more commonplace in occupational epidemiology studies in which investigators perform preliminary screening of multiple occupational exposures without a priori hypotheses. Interpretation of results from hierarchical regression analyses may mitigate some of the problems inherent in conventional analyses, by controlling for correlated exposures, enhancing the precision of estimates, and providing some adjustment for associations occurring by chance by incorporating prior knowledge into the analysis. As in any occupational study, collection of accurate exposure information is crucial, and exposure assessment for use in hierarchical modeling requires additional thought as to the comparability of information quality across multiple

exposures, and the potential impact of data quality for individual exposures on the panel of results.

ACKNOWLEDGMENTS

The authors thank Drs. Sander Greenland and John Witte for their helpful suggestions concerning the analysis and careful review of the manuscript.

REFERENCES

Greenland S. 1992. A semi-Bayes approach to the analysis of correlated multiple associations with an application to an occupational cancer-mortality study. Stat Med 11:219–230.

Greenland S. 1993. Methods for epidemiologic analyses of multiple exposures: a review and comparative study of maximum-likelihood, preliminary-testing, and empirical-Bayes regression. Stat Med 12: 717–736.

Greenland S. 1994. Hierarchical regression for epidemiologic analyses of multiple exposures. Environ Health Perspect 102:33–39.

Greenland S. 1997. Second-stage least squares versus penalized quasi-likelihood for fitting hierarchical models in epidemiologic analysis. Stat Med 16:515–526.

Greenland S. 1998. Introduction to regression modeling. In: Rothman K, Greenland S, editors. Modern epidemiology. Philadelphia, PA: Lippincott-Raven Publishers. p 401–432.

Greenland S. 2000. Principles of multilevel modelling. Int J Epidemiol 29:158–167.

Greenland S. 2000. When should epidemiologic regressions use random coefficients? Biometrics 56(3):915–921.

Greenland S, Poole C. 1994. Empirical-Bayes and semi-Bayes approaches to occupational and environmental hazard surveillance. Arch Environ Health 49:9–15.

Olshan AF, De Roos AJ, Teschke K, Neglia JP, Stram DO, Pollock BH, Castleberry RP. 1999. Neuroblastoma and parental occupation. Cancer Causes Control 10:539–549.

Witte JS, Greenland S. 1996. Simulation study of hierarchical regression. Stat Med 15:1161–1170.

Witte JS, Greenland S, Haile RW, Bird CL. 1994. Hierarchical regression analysis applied to a study of multiple dietary exposures and breast cancer. Epidemiology 5:612–621.

Witte JS, Greenland S, Kim LL. 1998. Software for hierarchical modeling of epidemiologic data. Epidemiology 9:563–566.

Witte JS, Greenland S, Kim LL, Arab L. 2000. Multilevel modeling in epidemiology with GLIMMIX. Epidemiology 11(6):684–688.